arXiv:1510.00431v1 [cond-mat.stat-mech] 1 Oct 2015

Towards an optimal flow: Density-of-states-informed replica-exchange simulations

Thomas Vogel* and Danny Perez†
Theoretical Division (T-1), Los Alamos National Laboratory, Los Alamos, NM 87545

Replica exchange (RE) is one of the most popular enhanced-sampling simulations technique in use today. Despite widespread successes, RE simulations can sometimes fail to converge in practical amounts of time, e.g., when sampling around phase transitions, or when a few hard-to-find configurations dominate the statistical averages. We introduce a generalized RE scheme, density-of-states-informed RE (g-RE), that addresses some of these challenges. The key feature of our approach is to inform the simulation with readily available, but commonly unused, information on the the density of states of the system as the RE simulation proceeds. This enables two improvements, namely, the introduction of resampling moves that actively move the system towards equilibrium, and the continual adaptation of the optimal temperature set. As a consequence of these two innovations, we show that the configuration flow in temperature space is optimized and that the overall convergence of RE simulations can be dramatically accelerated.

Sampling the phase space of Hamiltonians to estimate thermodynamic properties is one of the fundamental problems in statistical physics. However, direct approaches often dramatically fail due to the presence of large free-energy barriers between different regions of phase space. While many methods have been proposed to address this challenge, few have had an impact as significant as the replica exchange (RE) method [1–4]. Indeed, since its introduction RE has established itself as a workhorse in atomistic and coarse-grained simulations and is currently used to investigate a large variety of systems in many areas ranging from statistical physics [5–7], over biology and chemistry [8–10], to solid-state physics and materials science [11–13].

RE enhances the exploration of phase space by using a set of N individual simulations (often called "walkers") that evolve through Monte Carlo (MC) or molecular dynamics (MD) updates under different external parameters. For example, each walker might run at a different temperature T_i ($1 \leq i \leq N$), which is the situation we consider in the following (in this context, RE is often referred to as Parallel Tempering). After predefined time intervals τ_{RE}, the exchange of the current microstates between pairs of walkers is attempted and carried out with an appropriate probability; $W_{\mathrm{acc}}(U_i, U_j) = \min[1, \mathrm{e}^{\Delta\beta\,\Delta U}]$ in the case of canonical walkers, where $\beta_i = (k_{\mathrm{B}}T_i)^{-1}$ are the inverse temperatures of the heat baths and U_i the internal potential energies of the microstates. This exchange mechanism promotes configurational mixing by exposing replicas to external conditions (e.g., high temperatures) where free energy barriers can easily be overcome. It further provides a means for thermodynamic information to be transferred to conditions where the convergence of direct MC or MD simulations would require prohibitively long simulation times.

The exchange probabilities in RE strictly comply with detail balance, which ensures that the proper canonical distributions will be sampled at all temperatures. This property is extremely useful, because samples taken at each temperatures can be used without reweighting. It can however be restrictive, especially in the early stages of a simulation, when systems are far from equilibrium. This is related to the fact that conventional RE does not provide a natural mechanism to integrate and exploit information that becomes available during the simulation.

In this manuscript, we show how one such inexpensive and readily available source of information, namely, concurrent estimates of the density of states, $g(U)$, can significantly improve conventional RE. We leverage the (approximate) knowledge of $g(U)$ in two ways: first, we introduce a resampling operation akin to a Gibbs sampling move, which samples according to the $g(U)$-inferred canonical distribution over an ensemble of configurations previously visited by any replica; as we will show thereby explicitly steering the system toward equilibrium. Second, we use estimates of $g(U)$ to continuously improve the temperature set $\{T_i\}$. Because resampling breaks correlations along individual trajectories on each replica, the temperature set can be made optimal with respect to diffusion in temperature space [14].

The enabling factor of our approach is the concurrent estimation of the density of states $g(U)$. While it can be obtained by a number of techniques [15–18], we here rely on ideas from the Adaptive Biasing Force (ABF) formalism [19] for MD simulations. The key is to frame the problem as the estimation of the free energy $F_\beta(U) = -k_{\mathrm{B}}T \ln[g(U)\,\exp(-\beta U)]$ through its derivative s, which can be written in terms of microcanonical averages [19] as:

$$s(U) = \frac{\mathrm{d}F}{\mathrm{d}U} = -\left\langle \frac{\mathrm{d}}{\mathrm{d}t}(w \cdot p) \right\rangle_U , \qquad (1)$$

with p being the vector of momenta and $w = \nabla U/(\nabla U \cdot \nabla U)$. Here, the derivative with respect to time is understood as a derivative along a micro-canonical trajectory. In practice, $\mathrm{d}(w \cdot p)/\mathrm{d}t$ is measured periodically (say every 100 timesteps) and stored with its corresponding value of U. We then use binned averages to reconstruct $s(U)$. By integration, we recover $F_\beta(U)$, and hence $g(U)$ [20]. In an alternative representation, $s(U)$ can

be used to define a microcanonical temperature observable: $T_{\rm m}(U) = T_0/(1 - s(U))$, where T_0 is the heat-bath temperature. By integrating the thermodynamic relation $1/T_{\rm m}(U) = {\rm d}S(U)/{\rm d}U$, with $S(U) = k_{\rm B} \ln g(U)$, one reaches the same result. Note that if the momenta are unavailable, e.g., when using MC dynamics, a configurational temperature $T_{\rm m}(U)$ can be estimated based on structural data alone [21–23]. The cost of estimating $s(U)$ is negligible in practice. An advantage of this approach is that the validity of Eq. (1) does not depend on the sampling being carried out in any particular ensemble: the only requirement is that the dynamics yield equal probabilities of observing different configurations with the same U. To minimize the impact of initial — potentially far-from-equilibrium — states on the estimator at later times, we introduce a memory time (much larger than all other time scales) after which measurements are discarded.

While in the ABF method [19], $s(U)$ is used to create a multicanonical ensemble, we instead leverage it in g-RE in the following ways. While leaving the original RE mechanism untouched, we first introduce an additional, global resampling move (executed after time intervals $\tau_{\rm resamp}$) by which the microstate of a replica is resampled from the ensemble of configurations visited by *any* of its peers at *any* time in the past. In practice, this is enabled by a global configuration database populated by all walkers during the simulation. A configuration is selected from the database with a probability proportional to its estimated canonical weight $P_{\beta_i}(U) = g(U) \exp[-\beta_i U]$. This is akin to an approximate Gibbs sampling [24, 25]. What is crucial here is that the $P_{\beta_i}(U)$ are inferred from *global* thermodynamic information, so that it can differ from the distribution locally sampled by the corresponding walker. In other words, this operation actively steers the distributions towards what is globally deemed equilibrium. It does so by allowing for the *replication* of thermodynamically relevant states, in contrast to conventional RE where configuration can only *diffuse* in configuration space. As will be shown below, resampling proves essential for convergence in the neighborhood of strong phase transitions and for the timely escape out of metastable states. One might, however, wonder whether g-RE produces correct statistics, as resampling does not *a priori* obey global balance when the configuration database is finite. In fact, the only requirement for correctness is that any two state with the same U have the same probability of eventually being observed during an arbitrary long simulation.

As discussed in the Supplementary Materials below, this remains the case when resampling is introduced, as long as the samplers on each replicas (e.g., Langevin-MD or Metropolis MC) are ergodic and canonical in and of themselves. In that limit, $g(U)$, as obtained through Eq. 1 will converge to the correct value. From the knowledge of $g(U)$ and from a sample of observed configurations, any canonical quantity can then be obtained. It is however important to note that resampling introduces correlations between the configurations stored in the database at a specific point in time, and hence, potentially also between replicas. Our approach exploits these correlations to share information between replicas. However, resampling too frequently will lead to statistical ineficiencies. The optimal choice of $\tau_{\rm resamp}$ will be discussed in future works.

The availability of $g(U)$ also enables a second innovation: the continuous optimization of the temperature set $\{T_i\}$. The overall goal here is to minimize the round-trip time for replicas to wander between low and high temperatures. Much effort has been (and is still) dedicated to addressing this issue (see Refs. [14, 26–32] for examples). From this body of work it emerges that performance is often characterized in terms of two key concepts: the average exchange acceptance probabilities $W_{\rm acc}(i, i+1)$ between pairs of neighboring temperatures, and the flow ratio $f(i) = n_{\rm up}(i)/[n_{\rm up}(i) + n_{\rm down}(i)]$, i.e., the fraction of replicas that diffuse up in temperature for a given walker i (a replica is said to flow upwards if it visited the minimum temperature more recently than the maximum temperature, cf. [27, 28]). In the probability-centric view [14, 29, 31], the objective is to find the $\{T_i\}$ such that $W_{\rm acc}(i, i+1) = {\rm const}$, $\forall i$, the insight being that locally low acceptance probabilities would limit the free diffusion of replicas. In the flow-centric view [27, 28], the optimal set $\{T_i\}$ is such that $f(i) = 1 - [(i-1)/(N-1)]$ (for $T_i < T_{i+1}$, $\forall i$), as this indicates an unimpeded flow of replicas. In contrast, the presence of a bottleneck would signal itself by comparatively flat regions separated by a sharp drop in $f(i)$. The two approaches have been contrasted by Nadler and Hansmann [14] who showed that, in the special case where the dynamics on each replica become completely uncorrelated between exchange attempts, the *optimal* choice is to make the exchange probabilities constant, as this minimizes the round-trip time from the lowest to the highest temperatures. Furthermore, the flow ratio will also be linear in this case. In the current context, this limit can be approached by setting $\tau_{\rm resamp} = \tau_{\rm RE}$, as resampling decreases correlations (as long as τ_{RE} is not so short as to saturate the database with near-identical configurations). In g-RE, the optimal temperature set can easily be determined. Holding the minimal and the maximal temperatures fixed, we apply a bisection scheme until $\{T_i\}$ converges to a situation where

$$W^{\rm m}_{\rm acc}(P_{i+1}, P_i) = \int_{-\infty}^{\infty} P_{\beta_i}(U) \int_{-\infty}^{\infty} W_{\rm acc}(U, U') \, P_{\beta_{i+1}}(U') \, {\rm d}U {\rm d}U' = {\rm const}, \ \forall i \, . \tag{2}$$

(see Supplemental Materials below for more details). Here also, the $P_{\beta_i}(U)$ are based on the current estimator of $g(U)$, and Eq. (2) can be evaluated numerically. This

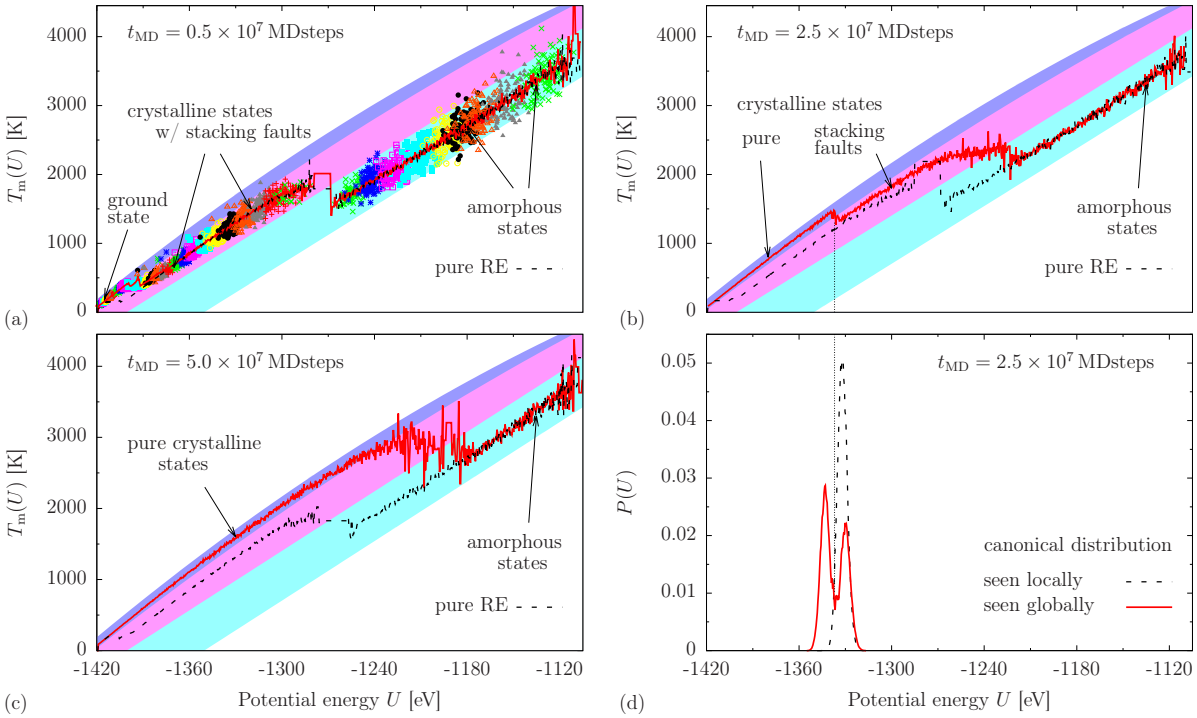

FIG. 1. Convergence of $T_m(U)$ at different simulation times: (a) shortly after the start of the run and after the first pure crystalline states have been discovered; (b) these states get replicated and move the spurious transition between pure and faulted states upwards in temperature; (c) after convergence of the g-RE scheme. The dots in (a) exemplarily show individual measurements, different colors correspond to different walkers (only a subset of data from every other walker is shown). The solid, red line shows $T_m(U)$, the dashed black line the same obtained from an equivalent, conventional RE run for comparison. The shaded bands in panels (a) to (c) are guides to the eyes and denote the type of configurations predominantly found in each of these regions, as indicated by the different labels. Panel (d) shows the estimated canonical distribution $P_{\beta_i}(U)$, at the time corresponding to panel (b), as obtained from global data (solid, red line) compared to what would be inferred by an isolated walker trapped in faulted states (dotted, black line).

scheme is very fast in practice and does not require any preliminary calculation. We continuously readjust the temperatures at intervals τ_{adapt}, assisting convergence in situations where the walkers begin far from equilibrium. Note that adaption of the $\{T_i\}$ does not interfere with the averages necessary for the measurement of $s(U)$, as samples taken at different temperatures can be seamlessly integrated through Eq. (1).

We now demonstrate the performance of our method on a system of 500 silver atoms. The $N = 71$ walkers are canonical molecular dynamics runs with a timestep of $\tau_{ts} = 2\,\text{fs}$ and the heat-bath temperatures are set by Langevin thermostats. Atoms interact via an embedded-atom potential [33]. The simulation cell is cubic and periodic boundary conditions are used in all directions. The minimum temperature is set to $T_{min} = 100\,\text{K}$ and the maximal to $T_{max} = 3500\,\text{K}$. The particle density is fixed at $\rho = 0.0585\,\text{Å}^{-3}$, which corresponds to the density that minimizes the energy of the fcc crystal; the fcc configuration is therefore the putative global energy minimum for this system and should dominate up to the vicinity of the melting point, given that the thermal concentration of defects is expected to be vanishingly small. However, all walkers are initialized from a quenched liquid (amorphous) configuration. This choice makes the system a very good prototype of a case where the thermodynamically relevant configurations are unknown *a priori* and difficult to access. Indeed, from MD studies of metals (e.g., Ag [34–36], Cu [37], Ni [38], etc.) it is known that recovering the perfect crystalline state from the melt requires very slow cooling, otherwise the system remains trapped in amorphous states (at fast cooling) or in a mixture of fcc and hcp regions separated by stacking defects (at moderate cooling rates). In addition, the presence of a first-order transition (melting) within the range of temperatures makes this an extremely challenging system to study. Finally, since the thermal distributions of crystal defects is expected to be vanishingly small in a system of that size away from the immediate vicinity of the melting point, the validity of the results is easy to assess.

As illustrated in Fig. 1, the coupling of the different replicas through the estimator of $g(U)$ and the configurational database directs the evolution of the system towards thermal equilibrium by enabling the *replication* of thermodynamically relevant, crystalline states, in contrast to conventional RE where states are only *exchanged*.

In that figure, we report the instantaneous estimators of the microcanonical temperature $T_{\rm m}(U)$ at different times in the simulation. In our example, the perfect crystalline states (upper, purple band) are essentially the only relevant ones below the melting point; they are, however, by construction not present in the early stages of the simulations. While, at lower temperatures, the system quickly leaves the amorphous region of phase space (lower, pale blue band) and crystallizes, most of these crystalline configurations initially contain stacking faults (middle, pink band) that take a very long time to anneal. Convergence of standard RE requires the perfect crystalline state to be *independently* found at least (and ideally much more than) $n_{<melt}$ times (the number of walkers for which $T_i < T_{\rm melt}$). The rate at which this occurs hence controls the convergence speed; from standard RE we infer this rate to be of the order $10^8\,{\rm s}^{-1}$ and the convergence would require approximately 2.5×10^8 MD steps. In contrast, resampling enables the replication of the pure crystalline state as soon at it is found *once* by *any* replica. As illustrated in Fig. 1d, this occurs because the *global* estimate of $g(U)$ eventually contains contributions from the crystalline region that are locally invisible to a walker trapped in faulted states. Computational gains follow because the probability of resampling a crystalline state significantly exceeds that of a replica independently finding it.

The same is true around the melting transition, which is initially biased to lower temperatures due to the initialization from a quenched liquid (note that the melting point at constant volume is much higher than the triple point). Before equilibration, a spurious transition between perfect crystals and faulted states (the sudden drop in $T_{\rm m}(U)$ at low energies) is observed and the melting temperature is underestimated. This is illustrated in Fig. 1 a–b (red, solid lines). A particular advantage of continuous temperature adaptation is that the temperature set remains nearly optimal at all times, even as the position of the (spurious or real) phase transitions evolve during convergence. With g-RE we observe convergence at $t_{\rm MD}\approx4\times10^7$ MD steps, while conventional RE, even with temperature adaptation (black, dashed lines), still mainly samples faulted states. The temperature adaptation scheme is very robust and converges quickly, even for a poor choice of the initial set $\{T_i\}$ (see below).

Finally, Fig. 2 shows how the introduction of global resampling at time scales comparable to the RE exchange time ($\tau_{\rm resamp}\approx\tau_{\rm RE}\sim10^2\tau_{\rm ts}$) and an adaptive temperature set results in an optimal flow of the replica through temperature space, i.e., *both* constant $W^{\rm m}_{\rm acc}(i,i+1)$ *and* linear $f(i)$ are observed (Fig. 2 d). The measured round-trip times ($\tau^{\rm m}_{\rm rt}\approx14000\,\tau_{\rm RE}$) are in perfect agreement with measured round-trip times for a purely random exchange process with corresponding exchange probabilities, which here correspond to the optimal limit [14]. As show in Fig. 2 a-c, these two conditions can only be obeyed when using both temperature adaptation and resampling. In this sense, our work integrates earlier efforts [27, 28, 31], where either constant exchange times, exchange rates, or the optimal flow had to be sacrificed.

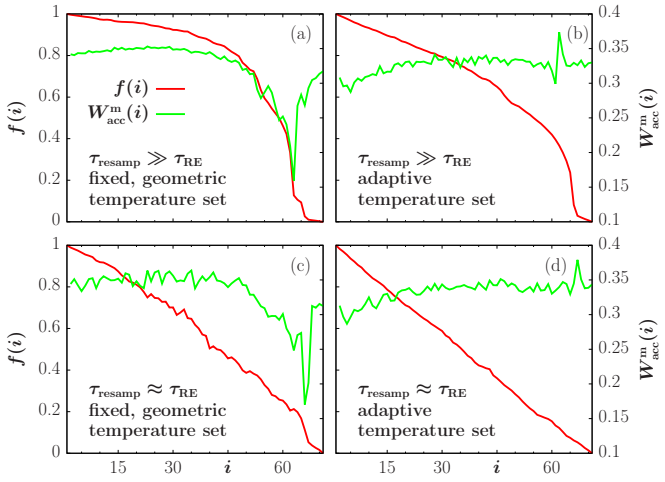

FIG. 2. Measured RE acceptance rates ($W^{\rm m}_{\rm acc}(i,i+1)$; green curves) and fraction of replicas diffusing from the lowest to the highest temperature ($f(i)$; red curves) for the 500 Ag atoms system. (a,c) fixed, geometric temperature set (b,d) adaptive temperature set; (a,b) non-ergodic sampler: $\tau_{\rm resamp}\gg\tau_{\rm RE}$, (c,d) ergodic sampler: $\tau_{\rm resamp}\approx\tau_{\rm RE}$.

In conclusion, we introduce a general scheme, g-RE, that can dramatically improve the efficiency of RE simulations. The method is based on the idea of informing RE simulations with estimators of the density of states $g(U)$ gathered on the fly. This allows for two key improvements: the introduction of a global resampling move that guides the system towards equilibrium and causes a dramatic reductions of correlation times of the sampling, and the on-the-fly determination of an optimal temperature set that simultaneously achieves constant exchange probabilities and linear flow ratio. We expect our method to be particularly useful for any system with dominant but hard to access states or around strong first-order transitions. So far, there have been two approaches to such problems: either to change the ensemble (see [29, 39] for examples) or to introduce global MC trial moves. In the former case, one often chooses to work in the multicanonical ensemble where $P(U)=$ const; creating such an ensemble is in fact a common way to leverage the knowledge of $g(U)$. This approach, however, is not optimal for the system reported here, because even though that ensemble is free from gradients of $F(U)$, the energy landscape locally remains extremely rough, thereby severely limiting the diffusivity in U. Using diffusion in T-space to promote mixing proved a more efficient alternative. Nonetheless, it still required the introduction of a global move to insure convergence. We here proposed a *generic* solution for constructing such a global move that does not require any *a priori* information about the system. We hence expect this approach to be useful for a

wide range of systems and to also be applicable to RE schemes using other ensembles [9, 12, 40–43].

This work was supported by the U.S. Department of Energy through the Los Alamos National Laboratory (LANL)/LDRD Program and used computing resources provided by the Los Alamos National Laboratory Institutional Computing Program. Los Alamos National Laboratory is operated by Los Alamos National Security, LLC, for the National Nuclear Security administration of the US DOE under contract DE-AC52-06NA25396. LA-UR-15-23605 assigned.

* tvogel@lanl.gov; Current address: Department of Physics, Stetson University, DeLand, FL 32723

† danny_perez@lanl.gov

[1] C. J. Geyer, in *Computing Science and Statistics: Proceedings of the 23rd Symposium on the Interface*, edited by E. M. Keramidas (Interface Foundation, Fairfax Station, VA, 1991) p. 156.
[2] A. P. Lyubartsev, A. A. Martsinovski, S. V. Shevkunov, and P. N. Vorontsov-Velyaminov, J. Chem. Phys. **96**, 1776 (1992).
[3] K. Hukushima and K. Nemoto, J. Phys. Soc. Japan **65**, 1604 (1996).
[4] D. J. Earl and M. W. Deem, Phys. Chem. Chem. Phys. **7**, 3910 (2005).
[5] M. Hasenbusch, A. Pelissetto, and E. Vicari, Phys. Rev. B **78**, 214205 (2008).
[6] B. Yucesoy, H. G. Katzgraber, and J. Machta, Phys. Rev. Lett. **109**, 177204 (2012).
[7] E. Bittner and W. Janke, EPL (Europhys. Lett.) **74**, 195 (2006).
[8] U. H. E. Hansmann, Chem. Phys. Lett. **281**, 140 (1997).
[9] J. Gross, T. Neuhaus, T. Vogel, and M. Bachmann, J. Chem. Phys. **138**, 074905 (2013).
[10] H.-H. G. Tsai, M. Reches, C.-J. Tsai, K. Gunasekaran, E. Gazit, and R. Nussinov, Proc. Natl. Acad. Sci. U.S.A. **102**, 8174 (2005).
[11] M. Falcioni and M. W. Deem, J. Chem. Phys. **110**, 1754 (1999).
[12] S. Auer and D. Frenkel, Nature **409**, 1020 (2001).
[13] F. Chuang, C. Ciobanu, C. Predescu, C. Wang, and K. Ho, Surf. Sci. **578**, 183 (2005).
[14] W. Nadler and U. H. E. Hansmann, Phys. Rev. E **75**, 026109 (2007).
[15] F. Wang and D. P. Landau, Phys. Rev. Lett. **86**, 2050 (2001).
[16] A. Laio and M. Parrinello, Proc. Natl. Acad. Sci. **99**, 12562 (2002).
[17] J. Kim, J. E. Straub, and T. Keyes, Phys. Rev. Lett. **97**, 050601 (2006).
[18] C. Junghans, D. Perez, and T. Vogel, J. Chem. Theory Comput. (JCTC) **10**, 1843 (2014).
[19] E. Darve, D. Rodríguez-Gómez, and A. Pohorille, J. Chem. Phys. **128**, 144120 (2008).
[20] For brevity, we will use the same symbol for the exact $g(U)$ and its estimator in the following. It should however be understood that, in the practical implementation of g-RE, the exact $g(U)$ is not available, so that $g(U)$ there refers to the estimator based on Eq. (1) instead.
[21] H. H. Rugh, Phys. Rev. Lett. **78**, 772 (1997).
[22] B. D. Butler, G. Ayton, O. G. Jepps, and D. J. Evans, J. Chem. Phys. **109**, 6519 (1998).
[23] See Supplemental Material at [URL will be inserted by publisher], which includes Refs. [44–47].
[24] J. S. Liu, *Monte Carlo strategies in scientific computing*, 2nd ed. (Springer-Verlag, New York, NY, 2002).
[25] J. D. Chodera and M. R. Shirts, J. Chem. Phys. **135**, 194110 (2011).
[26] N. Rathore, M. Chopra, and J. J. de Pablo, J. Chem. Phys. **122**, 024111 (2005).
[27] H. G. Katzgraber, S. Trebst, D. A. Huse, and M. Troyer, J. Stat. Mech.: Theory Exp. **2006**, P03018 (2006).
[28] S. Trebst, M. Troyer, and U. H. E. Hansmann, J. Chem. Phys. **124**, 174903 (2006).
[29] T. Neuhaus, M. P. Magiera, and U. H. E. Hansmann, Phys. Rev. E **76**, 045701 (2007).
[30] A. Patriksson and D. van der Spoel, Phys. Chem. Chem. Phys. **10**, 2073 (2008).
[31] E. Bittner, A. Nußbaumer, and W. Janke, Phys. Rev. Lett. **101**, 130603 (2008).
[32] A. J. Ballard and D. J. Wales, J. Chem. Theory Comput. (JCTC) **10**, 5599 (2014).
[33] P. L. Williams, Y. Mishin, and J. C. Hamilton, Model. Simul. Mater. Sci. Eng. **14**, 817 (2006).
[34] S. Xiao and W. Hu, J. Chem. Phys. **125**, 014503 (2006).
[35] Z.-A. Tian, R.-S. Liu, H.-R. Liu, C.-X. Zheng, Z.-Y. Hou, and P. Peng, J. Non-Cryst. Solids **354**, 3705 (2008).
[36] Z. Jian, J. Chen, F. Chang, Z. Zeng, T. He, and W. Jie, Sci. China Technol. Sci. **53**, 3203 (2010).
[37] C. S. Liu, J. Xia, Z. G. Zhu, and D. Y. Sun, J. Chem. Phys. **114**, 7506 (2001).
[38] F. Li, X. Liu, H. Hou, G. Chen, and G. Chen, Intermetallics **19**, 630 (2011).
[39] J. Kim, T. Keyes, and J. E. Straub, J. Chem. Phys. **132**, 224107 (2010).
[40] Y. Sugita and Y. Okamoto, Chem. Phys. Lett. **329**, 261 (2000).
[41] J. Kim, J. E. Straub, and T. Keyes, J. Phys. Chem. B **116**, 8646 (2012).
[42] T. Vogel, Y. W. Li, T. Wüst, and D. P. Landau, Phys. Rev. Lett. **110**, 210603 (2013).
[43] T. Vogel, Y. W. Li, T. Wüst, and D. P. Landau, Phys. Rev. E **90**, 023302 (2014).
[44] M. P. Allen and D. Quigley, Mol. Phys. **111**, 3442 (2013).
[45] C. Junghans, M. Bachmann, and W. Janke, Phys. Rev. Lett. **97**, 218103 (2006).
[46] S. Schnabel, D. T. Seaton, D. P. Landau, and M. Bachmann, Phys. Rev. E **84**, 011127 (2011).
[47] R. L. Burden and J. D. Faires, *Numerical Analysis*, 3rd ed. (PWS Publishers, 1985).

SUPPLEMENTAL MATERIAL

Microcanonical temperature measurements

In this work we rely on methods that use measurements of temperature-like observables in order to estimate the derivative of the entropy (or, equivalently, of the density of states) with respect to energy. For example, the

microcanonical temperature can be written in terms of configurational properties as

$$\frac{1}{k_B T_{\mathrm{conf}}(U)} = \left\langle \nabla \cdot \left(\frac{\nabla U}{\nabla U \cdot \nabla U} \right) \right\rangle_U \tag{3}$$

where $\langle \ldots \rangle_U$ is the average over states with a given potential energy U [21, 22]. For canonical averages or for very large system sizes, $T_{\mathrm{conf}}(U)$, measured from configurations in a heat-bath at constant temperature, is equal to just that canonical temperature [22, 44]. Another way of estimating $\partial S/\partial U$ is based on thermodynamic integration using time derivatives, as derived in the context of the adaptive biasing force (ABF) method [19]. By using the potential energy as the 'order parameter', $T_{\mathrm{m}}(U) = (\partial S(U)/\partial U)^{-1}$ can be estimated via

$$T_{\mathrm{ABF}}(U) = T_0/(1 - s(U))\,, \tag{4}$$

where T_0 is the heat-bath temperature of the particular canonical-ensemble MD walker and with

$$s(U) = -\left\langle \frac{\mathrm{d}}{\mathrm{d}t} \left(\frac{\nabla U}{\nabla U \cdot \nabla U} \right) \cdot p \right\rangle_U , \tag{5}$$

where p are the particle momenta. In contrast to Eq. (3), we do not need to calculate second derivatives of the potential energy, but only a time-derivative obtained in finite-difference form using a constant-energy MD timestep. This can be much less costly, especially for a complex many-body potentials, such as the one used in our work. However, if p are not available, in a MC simulation, for example, one can instead use $T_{\mathrm{conf}}(U)$. In Fig. 3 we show raw measurements T' that contribute to the averages in Eqs. (3) and (4). Every data point corresponds to measurements of T'_{ABF} and T'_{conf} for one particular microstate sampled during the run. An extremely good agreement can be seen for both observables, not only with respect to each other but also with respect to the canonical reference temperatures indicated by the dotted lines. After carrying out the averages, we do not see any difference in the results of our simulations when using either of the observables and make our choice solely based on the computational performance.

As a valuable side-product, note that since $T_{\mathrm{m}}(U)$ is *measured*, the results are readily available for a microcanonical analysis of finite-size phase transitions [45, 46].

Determination of temperature sets

The transition (exchange) probability between two energy distributions P_{i+1} and P_i corresponding to different temperatures β_i and β_{i+1} takes the form

$$W^{\mathrm{tot}}_{\mathrm{acc}}(P_{i+1}, P_i) = \iint_{-\infty}^{\infty} P_i(U) W_{\mathrm{acc}}(U, U') P_{i+1}(U')\, \mathrm{d}U \mathrm{d}U'. \tag{6}$$

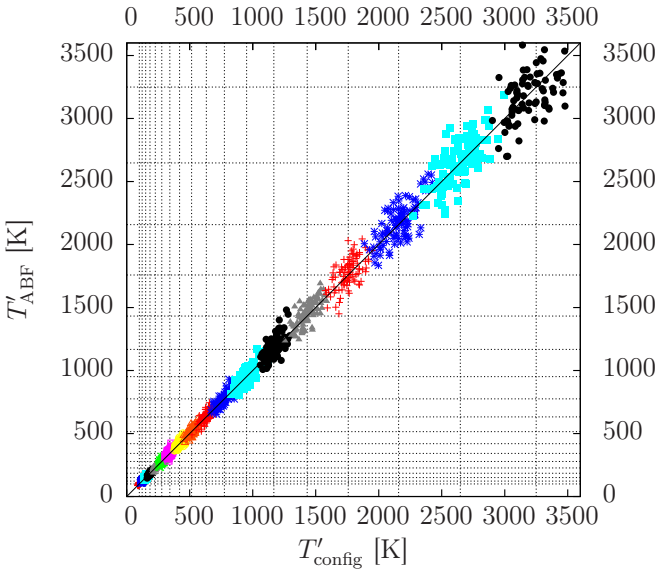

FIG. 3. T_{config} vs. T_{ABF} for one example g-RE run with a fixed temperature set for the $N = 500$ Ag system. Black, diagonal line: $T_{\mathrm{config}} = T_{\mathrm{ABF}}$ to guide the eye; Dotted lines: individual reference temperatures T_i. See text for details.

$W_{\mathrm{acc}}(U, U')$ is the usual acceptance probability for a replica exchange in canonical parallel tempering simulations: $W_{\mathrm{acc}}(U, U') = \min[1, \mathrm{e}^{\Delta\beta\Delta U}]$, with $\Delta\beta$ being the difference in inverse temperature and ΔU the energy difference of the corresponding replicas. The integral in Eq. (6) is evaluated numerically as a sum over histogram bins. Recalling that $g(U)$ is used to calculate the energy distributions $P_i(U) = g(U) \exp[-\beta_i U]$, it is clear that the integrand takes non-negligible values only in a limited energy range, and hence that the bounds of the integrals can be restricted in practice. For example, the estimator of $g(U)$ is zero (by construction) at energies lower than the lowest energy observed by any replica, hence so is $P_i(U) = g(U) \exp[-\beta_i U]$. Similarly, at high energies, the Boltzmann factor leads to a rapid decay of $P_i(U)$, so that the integral can be truncated when $P_i(U)$ becomes negligible.

To find the optimal temperature set, we first (arbitrarily) select a target exchange probability c_t and set $\beta_1 = 1/k_B T_{\mathrm{min}}$. We then successively apply a binary search (also known as interval halving method or bisection method; a standard method in numerical analysis [47]) to neighboring pairs of walkers to find β_{i+1} such that $W^{\mathrm{tot}}_{\mathrm{acc}}(P_{i+1}, P_i) = c_t$. In practice, we first bracket the interval of β_{i+1} that contains the solution to $W^{\mathrm{tot}}_{\mathrm{acc}}(P_{i+1}, P_i) = c_t$, and then apply a conventional binary search to identify the precise value. Given that the number N of replicas is constant, a given value of c_t implies a given value of T_N. To ensure that we always cover the same temperature range, we then apply an additional bisection scheme where we determine c_t

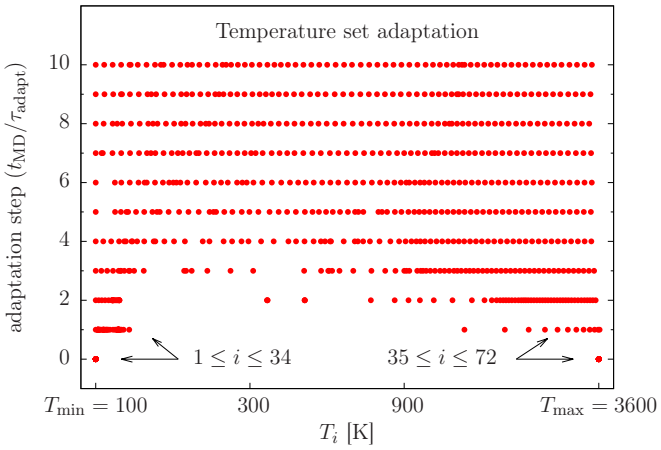

FIG. 4. Convergence of the temperature set $\{T_i\}$ after an initialization where every walker starts either at the lower or the upper end of the total temperature range. The plot shows the evolution of $\{T_i\}$ during the first ten adaptation steps; after very few steps, the simulation optimally covers the whole temperature range. Data measured during the same simulation run shown in Fig. 1 in the main article.

such that T_N falls within a small range around the targeted maximal temperature T_{max}. That is, whenever T_N turns out to be too small we decrease c_t (and vice versa), in ever smaller amounts [47] and repeat the procedure described above. Of course, if the converged c_t is too small, one should increase the number of replicas N. Note that Eq. (2) (in the main article) only needs to be satisfied to some moderate precision to reap most of the performance benefits. Therefore, the tolerance in the bisection scheme, i.e., the range around the target value which is acceptable, can be adjusted. In all our simulations we observed that the determination of the optimal temperature set required only a negligible amount of computing time compared to the cost of actual MD, even though one evaluates a double-integral within a double-bisection scheme.

To emphasize the robustness of our scheme, we choose a far-from-optimal initial set $\{T_i\}$ for the run shown in Fig. 1 in the main article. About half of the walkers start at T_{min}, the other half at T_{max}; where there is, for all practical purposes, no overlap at all between the corresponding canonical distributions. Furthermore, recall that all walkers start with the same initial configuration. Still, the temperature set spreads out rapidly (see Fig. 4, in analogy to corresponding figures in Refs. [27, 28]) so that the walkers cover the whole range of temperatures and energies after a few adaptation steps. This illustrates that the method does not depend on the availability of any prior information on the system – in particular, no assumption or pre-estimation of the temperature set is necessary – and that it is very robust against even extreme choices of initial conditions.

Convergence to the canonical distribution

Despite of the fact that the g-RE procedure is history-dependent, i.e., that the result of the resampling depends on the previous history of the simulation through the content of the database, the method nonetheless converges to the canonical distribution. This can be seen by first considering a simpler situation: a single walker evolving through a canonical sampler and a database resampler. In this case, our database resampling procedure becomes:

- Sample a new configuration from the database
- Propagate for some time Δt with a canonical sampler
- Insert the current configuration back into the database
- If the number of configuration in the database exceeds M, remove an element at random. Equivalently, old configurations could be purged from the database after a set time (as done in the manuscript).
- Repeat

Following this prescription, after N cycles of the algorithm, every member of the database is the product of a proper time evolution. Indeed, from the perspective of individual configurations, they are randomly sampled, propagated for some time t with a canonical sampler, and then replaced in the database, i.e., *they are the product of an unbiased sequence of propagation with the correct dynamics*. However, because the configurations in the database are correlated, resampling from the *same* database does not produce independent canonical samples. For example, consider the case $M = 2$: both configurations are now likely to be extremely similar if the propagation time Δt is short. As we do not require independence, this is not an issue for g-RE in practice.

The second crucial point is that the "age" of the configurations in the database, i.e., the total (aggregate) amount of time they were propagated for, scales with the total simulation time. Therefore, in spite of correlations between individual entries, the sampling from the database over long times will be correct. Introducing a memory time after which configurations are dropped, as we did in the manuscript, directly imposes such a condition: "young" configurations cannot persist in the database by construction. Numerical simulations of the random deletion procedure demonstrate that this is also true there, to a remarkable extend. In fact, the mean age of a configuration is $\sim N\Delta t/M$, as can be intuitively expected. What is perhaps surprising is that the standard deviation of the age is a time (N)-independent constant that scales as $M^{1/2}$, i.e., the relative fluctuation (standard deviation of the age over the average age) *vanishes*

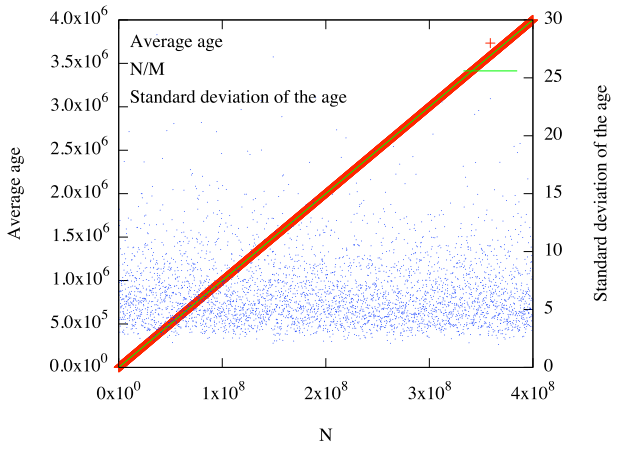

FIG. 5. Statistics of the age of configurations in the database vs. N for $M = 100$.

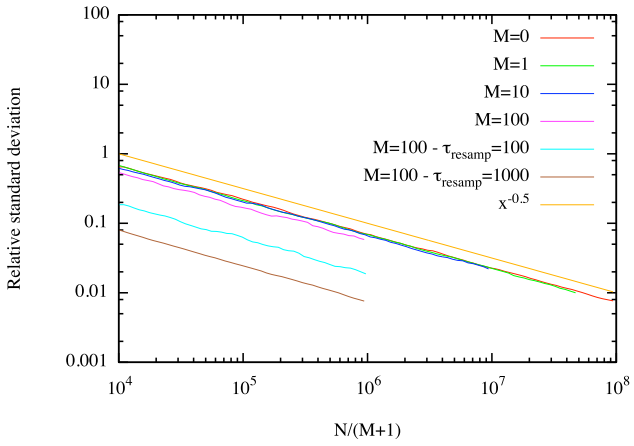

FIG. 6. Relative standard deviation (w.r.t. the average) of the position x of a walker for various memory size and resampling times. Unless otherwise specified, the resampling time is 1.

at large N. Results for a memory size of 100 are shown in Fig. 5.

This shows that, while the average age of each configuration in the database is smaller than N, it scales with simulation time, and it is exceedingly unlikely that a very "young" configuration will be sampled. Taken together with the fact that each one of them is the product of a proper time evolution, this shows the above procedure yield a canonical sampler over long enough times. In the worst case (when Δt is small) the configurations in the database will form a tight bundle in phase space. However, the bundle itself will evolve according to proper dynamics. Therefore, sampling from the database will yield unbiased canonical results over long times, so that we have an equal probability of sampling from the set of states with the same potential energy. As discussed in the manuscript, this is the only requirement for g-RE to be a valid sampler. However one should not expect multiple samples obtained from the database *at the same time* to be independent. Note that the same argument translates directly to the RE case, i.e., each configuration in the database taken alone, is a the product of proper dynamics whose age scales with the simulation time. We note that the situation appears problematic in the infinite M limit or in the infinite memory time limit. As this is not a situation we have to confront in practice (because of finite computer memory), we do not explore that limit further.

The above discussion has been demonstrated using direct numerical simulations. We now consider a two-replica case. For simplicity, the system is a 1D square well in $0 < x < 1$. Given the flat energy landscape, the temperature is irrelevant, so doing RE would yield acceptance with unit probability and thus would yield identical results. The dynamics are carried out with Metropolis MC with a Gaussian proposal distribution of standard deviation 0.01. Unless otherwise noted, we carry out only one MC step between database resampling ($\tau_{\rm resamp} = 1$). In Fig. 6, we report the relative standard deviation (the standard deviation over the average) of an histogram of the visited positions of one of the walkers as a function of $N/(M+1)$ for $M = 0$ (standard MC), 1, 10, and 100. Note that each walker populates its own database but can only sample from the database of the other walker (which is equivalent to the setup used in the manuscript). Since the canonical distribution is uniform, convergence is indicated by a vanishing relative deviation at long times. As shown in the figure, all cases converge at roughly the same rate once the average age in the database is accounted for. Note that, because of decreased correlations due to the larger database, the $M = 100$ case is converging even faster than expected from simple scaling arguments. Indeed, running the same case but resampling only every 100 or 1000 MC steps, thereby further limiting the correlation between entries in the database, yields an even faster convergence. Eventually, in the limit where the resampling time exceeds the ergodic relaxation time of the dynamics, the standard MC results (vs. N, not N/M as plotted) are recovered, which is the expected behavior.